Using Knowledge Maintenance for Preference Assessment

Nilesh L. Jain, DSc, and Michael G. Kahn, MD, PhD
Section of Medical Informatics, Department of Internal Medicine;
Department of Computer Science; Washington University, St. Louis, Missouri

Most real-life decisions require the decision maker to make trade-offs in order to fulfill multiple conflicting objectives. This is especially true in medical decision making while selecting the optimal therapy plan from among competing therapy plans for a patient. Multiattribute utility theory provides a framework to specify these trade-offs for optimal decision making based on the preferences of the decision maker. However traditional preference-assessment techniques are difficult to implement and rarely elicit the true preferences of the decision maker. We describe a new preference-assessment method based on the concept of knowledge maintenance where the preference model is changed each time it makes an incorrect recommendation. The method is implemented in a decision-theoretic system to evaluate competing three-dimensional radiation treatment plans. The preference-assessment method leads to preference models which perform better than preference models elicited using traditional assessment techniques.

INTRODUCTION

Computer-based clinical decision-support systems can be divided into two broad categories—heuristic knowledge-based systems and decision-theoretic preference-based systems. Most early research in medical informatics on clinical decision-support systems focused on using artificial intelligence techniques for building medical expert systems.[1,2] Recently there is growing interest in constructing decision-theoretic decision-support systems.[3,4] It is obvious that the performance of these systems depends critically on the underlying knowledge bases and preference models.

One of the main problems facing the developers of these systems is the elicitation of information to be contained in the systems—knowledge acquisition for knowledge-based systems, and preference assessment for preference-based systems. While computer scientists have focused on knowledge acquisition, much less effort has been invested in preference assessment. This paper describes a preference-assessment technique for one class of decision-theoretic systems—systems which objectively evaluate alternate decisions and select the best decision to fulfill multiple conflicting objectives. This includes the class of medical decision-making problems that involves the selection of the optimal therapy plan from competing therapy plans for a patient. We applied this technique to the clinical decision problem of ranking competing three-dimensional (3D) conformal radiation treatment (CRT) plans.

We first present some background information on decision theory, knowledge acquisition, preference assessment, and radiation treatment planning. We then describe a decision-theoretic model for ranking competing 3D CRT plans. We then describe our new preference-assessment method and an experiment of its implementation. We analyze the results of our experiment and make some concluding remarks.

BACKGROUND

Decision Theory

Decision theory provides a framework for selecting the optimal decision from many decision alternatives.[5] The two cornerstones of decision theory are its ability to handle uncertainty and its ability to handle multiple conflicting objectives.

We focus on decision problems which require the decision maker to make trade-offs in order to fulfill multiple conflicting objectives, as is the case in most real-life decisions. Multiattribute utility theory provides a framework to specify these trade-offs and to select the optimal decision according to the preferences of the decision maker.[6,7] The decision problem is decomposed into a number of meaningful attributes corresponding to the multiple objectives. The utility of each attribute is assessed and indicates how closely the attribute achieves its objective. Weights signifying the trade-offs among the attributes are acquired. Multiattribute utility theory provides various combining functions in order to obtain the overall utility for each decision, and the decision with the maximum overall utility is chosen.

Clinical decision analysis has been well studied and shown to address the problems faced in routine medical decision making.[8] We focus on the selection of the optimal therapy plan from competing therapy plans.

0195–4210/95/$5.00 © 1995 AMIA, Inc.

Knowledge Acquisition

Knowledge acquisition occurs through the interaction between the knowledge engineer developing a knowledge-based system and the domain expert whose expertise will be represented in the system.[9] It is a major focus of current research on knowledge-based systems.[10] Researchers have developed innovative knowledge-acquisition techniques for different classes of knowledge-based systems.[11] Knowledge maintenance is the addition, deletion, or modification of knowledge in a knowledge base.[12]

Preference Assessment

Preference assessment occurs through the interaction between the decision analyst developing the decision-theoretic system and the decision maker whose preferences will be represented in the system.[13] It is a major focus of current research on decision-theoretic systems.[14] There are two categories of preference-assessment techniques—for assessing attribute utilities and for assessing attribute weights.

Decision theory provides various standard techniques for both categories. However, multiple studies have shown the difficulty in actually implementing these techniques or getting accurate preferences.[15,16] We focus on the assessment of trade-off weights. The chief drawback of the standard weight-assessment techniques is that they ask the decision maker to consider hypothetical situations which never occur in reality, making it difficult for the decision maker to express a true preference.[17,18,19,20]

Recent Techniques

Two recent techniques proposed by medical informatics researchers—one for knowledge acquisition, and one for preference assesssment—provided the ideas for our preference-assessment technique. We describe them briefly here.

Ripple Down Rules. This is a knowledge acquisition technique based on knowledge maintenance.[21] It was first implemented in PEIRS (Pathology Expert Interpretative Reporting System), a rule-based expert system for providing clinical interpretations of data from thyroid tests.[22] Whenever the expert system makes an inadequate or incorrect recommendation, the expert adds or modifies a rule in the knowledge base. The key idea provided by this technique which we have applied to preference-based systems is that preference assessment can be done by updating the preference model whenever the recommendations of the system and the decision maker do not agree.

Simulated Decision Scenarios. This is a preference-assessment technique which infers the preference model of the decision maker based on her decisions on simulated decision problems.[23] This approach was implemented in VentPlan which calculates the recommended settings for four controls of a ventilator by evaluating the predicted effects of alternative ventilator settings.[24] The key idea provided by this technique is that preference assessment can be done during the actual use of the decision-support system.

Radiation Treatment Planning

Effective radiation treatment involves uniform irradiation of all tumor volumes to their prescribed high doses, while minimizing irradiation to nearby normal tissues.[25] Thus in the evaluation and selection of a radiation treatment plan, the radiation oncologist must make trade-offs in the radiation doses delivered to tumor volumes and normal tissues. This evaluation was easy in traditional two-dimensional treatment planning using coplanar standard radiation beam configurations because it involved comparing doses in only one or a few planes.

State-of-the-art treatment planning techniques use multiple non-coplanar radiation beams to design individualized 3D conformal plans where the shape of the high prescription-dose volume is conformed to the shape of the tumor, minimizing radiation to the nearby normal tissues.[26] Evaluation of these plans is difficult because it requires the radiation oncologist to decipher a huge amount of numerical and graphical data. We are using multiattribute utility theory in order to develop an objective plan-evaluation model for competing 3D CRT plans.[27]

OBJECTIVE PLAN-EVALUATION MODEL

Each attribute represents a clinical issue such as non-eradication of the tumor or radiation-induced damage to a nearby normal tissue. Each attribute has a *utility* (from 0 to 1) to indicate how closely the objective for that attribute is achieved, and a *weight* (from 0 to 1) to make trade-offs among the different attributes. The contribution of each attribute to the overall utility is called its *score*. When *utility* is low and *weight* is high, *score* should be low. When *utility* is high or *weight* is low, *score* should be high. One function which has this behavior is:

$$score_i = 1 - (1 - utility_i) \times weight_i \qquad (1)$$

The overall utility of the plan is called its *figure of merit (FOM)*. When *score* for any attribute is low, *FOM* should be low. Since *score* is between 0 and 1, the multiplicative model is a suitable combining function. Thus, we have:

$$FOM = \prod_{i}^{attr} (1 - (1 - utility_i) \times weight_i) \quad (2)$$

The two building blocks of the model are *utility* and *weight*. We describe each in further detail. The preference model of a radiation oncologist for a tumor site is the set of his weights for the attributes.

UTILITY MODEL

The *utility* of an attribute indicates how well its objective is met. For tumor volumes, the objective is to irradiate them to their prescription dose. For normal tissues, the objective is to minimize the dose to them. Thus *utility* has to be a function of the 3D dose distribution in the tissue represented by that attribute. However, it is impractical if not impossible to enumerate all the possible 3D dose distributions for a tissue. This makes it impossible to elicit utility functions based on dose distributions.

Instead, we use *proxy attributes* to elicit the utility functions. A proxy attribute reflects the degree to which an associated objective is met but does not directly measure the objective.[6] We experimented with a series of proxy attributes before arriving at one which was suitable. The proxies we considered but rejected were:

- probability of a bad outcome;
- percent of outlined tissue volume above its prescription or tolerance dose (called % volume);
- single dose statistic such as minimum, maximum, or mean dose.

We found each of them to be inadequate individually. The proxy that we use is a set of multiple dose statistics. For each attribute, the radiation oncologist is asked to choose one or more of the following dose statistics in the order of importance for comparing a tissue in different plans:

- minimum or maximum dose;
- mean dose;
- % volume.

To maximize tumor irradiation, the minimum dose, mean dose, or % volume should be as high as possible. To minimize normal tissue irradiation, the maximum dose, mean dose, or % volume should be as low as possible. Thus each of these dose statistics is suitable as a proxy attribute.

For each attribute, we obtained a set of dose statistics in the order of importance. The selection of the appropriate proxy dose statistic for an attribute was done as follows. Given two plans and a tissue, the model first compared the values of the most important dose statistic for that tissue. If the lower value was within 5% of the higher value, the model compared the next dose statistic. This continued till the model arrived at the first dose statistic which differed by at least 5%. If none of the dose statistics differed by 5%, then the model chose the last dose statistic as the two dose distributions were nearly equivalent. The choice of 5% was arbitrary and was suggested by the radiation oncologists.

We obtained utility curves for each dose statistic and tissue combination by direct elicitation. We obtained dose statistic values for extreme utilities of 0 and 1, and elicited a few intermediate points to interpolate a utility curve. Table 1 contains the attributes for prostate tumors with dose statistics in the order of importance specified by the radiation oncologist.

Table 1: Prostate tumor attributes and dose statistics.

Attributes	Dose statistics
Tumor volume	% volume, minimum, mean
Rectum, Bladder Femoral heads	% volume, maximum, mean
Connective tissue	maximum, % volume

We also obtained the attributes, importance-ordered dose statistics, and utility curves for lung and abdominal tumors.

INITIAL PREFERENCE MODEL

We elicited the initial set of weights for a tumor site using a variant of the trade-off technique.[28] The radiation oncologist selected the most important attribute i_c and it was assigned a weight of 1. For every other attribute i, we created two hypothetical plans P1 and P2. Table 2 contains the dose statistic values for the attributes. A dose statistic value of *ds0* means that *utility* is 0, and *ds1* means that *utility* is 1.

Table 2: Dose statistic values for hypothetical plans P1 and P2 to elicit initial weight of attribute i.

	Attribute i_c	Attribute i	Other attributes
Plan P1	*ds1*	*ds0*	*ds1*
Plan P2	*ds0*	*ds1*	*ds1*

We asked the radiation oncologist to select the preferred plan. By assigning a *utility* of 1 to all other attributes, the radiation oncologist is asked to trade off between attributes i_c and i. Three cases were possible:

1. Plans P1 and P2 were equivalent. Since the two attributes had complementary utilities, they must have the same weight. So $weight_i = 1$.
2. Plan P1 was preferred. In this case, we asked the radiation oncologist to give a new dose statistic value ds for attribute i_c in plan P2 to make the two plans equivalent. Then, by equating the *FOM* values using Eq. 2, we get $weight_i = 1 - U(ds)$ where U is the utility function of i_c.
3. Plan P2 was preferred. This was not possible as it indicated that i was more important than i_c.

We used this method to get the initial preference model for prostate, lung, and abdominal tumors from three radiation oncologists—one for each tumor site. Table 3 contains the initial preference model for prostate tumors.

Table 3: Initial preference model.

Attribute	Weight
Tumor volume	1.00
Rectum	0.17
Bladder	0.31
Femoral heads	0.91
Connective tissue	0.61

NEW PREFERENCE-ASSESSMENT METHOD

The previous section describes the traditional method for assessing preference models of radiation oncologists. The clinical experts found it conceptually difficult because the hypothetical plans P1 and P2 never occur in actual clinical practice. So we assume that the preference model assessed is only approximately correct, and will lead to some incorrect rankings. We call it the initial preference model. We use the concept of knowledge maintenance for refining the preference model, similar to the work on Ripple Down Rules described earlier. Each time the decision-support system makes a recommendation that does not agree with the expert decision maker (radiation oncologist), we ask him to examine the preference model to find the reason for the inconsistency and modify the preference model. The decision-support system always uses the current preference model of the decision maker. After sufficient use, we hypothesize that the preference model will converge to the correct set of weights resulting in the final preference model.

The preference-assessment method assumes that no gold standard is available for the decision problem. Thus it attempts to elicit the decision-making capability of an expert decision maker, similar to knowledge acquisition for knowledge-based systems. It does so by using the subjective decision-making of the expert decision maker as the gold standard, and making him provide his preferences for an objective decision model. The method is useful for decision problems which represent multiple similar decision-making opportunities in order to have sufficient number of cases for the preference assessment.

PREFERENCE-ASSESSMENT EXPERIMENT

We performed the preference-assessment experiment on three tumor sites—prostate, lung, and abdomen. We describe the experiment and results for prostate tumors. Data from 23 patients with stage A or B prostate cancer treated with 3D CRT at the Radiation Oncology Center of the Mallinckrodt Institute of Radiology in the last two years were available. We selected two different treatment plans for each patient. The radiation oncologist first ranked these two plans, and we called it the subjective ranking. Then we used the objective plan-evaluation model to obtain the objective ranking.

Rankings were obtained using an interactive decision-support system that allowed the radiation oncologists to examine the dose statistics, utilities, and weights of all the attributes, and modify the weights. The system was linked to our clinical treatment planning system[29] to have access to patient data. The system provides instantaneous feedback whenever the preference model is changed to assist in preference assessment. The graphical user interface of the system allowed the radiation oncologists to use the system unassisted.

If the subjective and objective rankings agreed, the radiation oncologist proceeded to the next patient case. If they did not, then the radiation oncologist examined and changed the weights appropriately until the rankings agreed. All changes in weights were logged. To make the assessment problem tractable, we assumed that the utility curves were correct.

Table 3 contains the initial preference model for prostate tumors. Figure 1 graphically shows the variation in weights of all five attributes. The weight for planning target volume changed once—from 1.00 to 0.85 while evaluating the plans for patient 10. The weight for rectum changed twice—from 0.17 to 0.80 while evaluating the plans for patient 1, and from 0.80 to

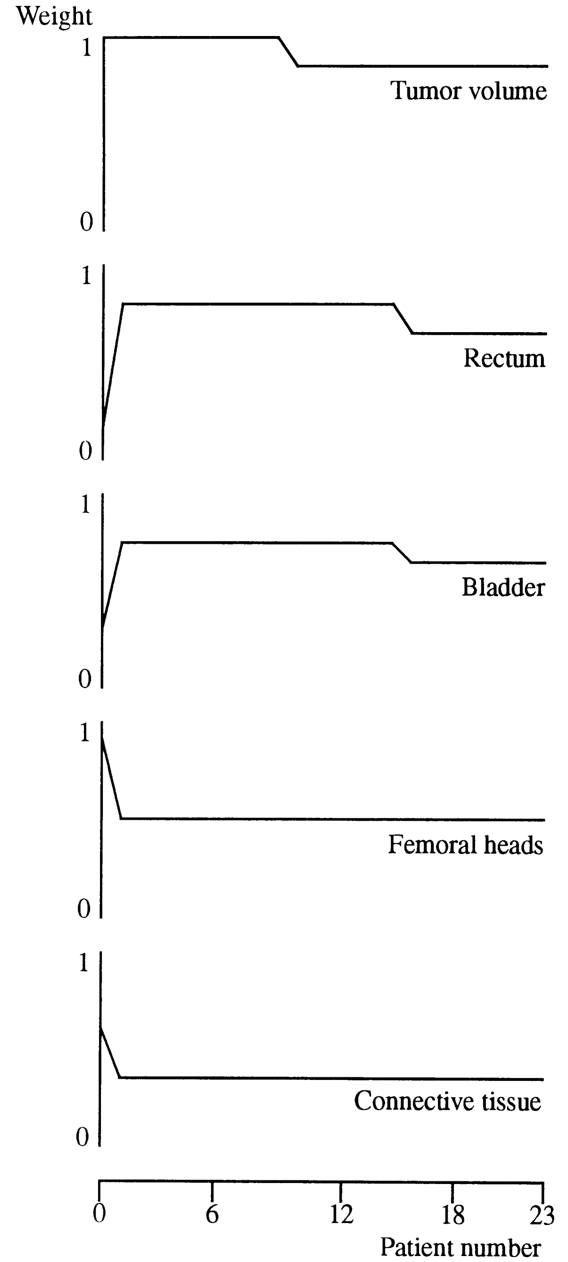

Figure 1: Variation in weight of the various attributes for prostate tumors.

0.65 while evaluating the plans for patient 16. The weight for bladder changed twice—from 0.31 to 0.75 while evaluating the plans for patient 1, and from 0.75 to 0.65 while evaluating the plans for patient 16. The weights for the femoral heads changed once—from 0.91 to 0.50 while evaluating the plans for patient 1. The weight for connective tissue changed once—from 0.61 to 0.35 while evaluating the plans for patient 1.

The preference model was changed three times—while evaluating the plans for patients 1, 10, and 16. The set of weights up to the first modification is the initial preference model. The set of weights between the first and second modification is called the first intermediate preference model. The set of weights between the second and third modification is called the second intermediate preference model. And the set of weights after the third modification is called the final preference model (see Table 4).

Table 4: Intermediate and final preference models.

Attribute	Preference model		
	First int.	Second int.	Final
Tumor volume	1.00	0.85	0.85
Rectum	0.80	0.80	0.65
Bladder	0.75	0.75	0.65
Femoral heads	0.50	0.50	0.50
Connective tissue	0.35	0.35	0.35

We observed similar variations in weights in lung and abdominal tumors. Thus we used the preference-assessment method with three different radiation oncologists for three different tumor sites.

ANALYSIS OF PREFERENCE MODELS

We have examined the performance of the four preference models (Tables 3 and 4). We did this by computing the agreement rate for each preference model. The agreement rate for a preference model is the fraction of patients for which the subjective ranking given by the radiation oncologist agrees with the objective ranking of the plan-evaluation model using that preference model. Table 5 contains the agreement rates for the four preference models.

Table 5: Agreement rates of preference models.

Preference Model	Agreement Rate
Initial	18/23 (78.3%)
First intermediate	21/23 (91.3%)
Second intermediate	22/23 (95.7%)
Final	22/23 (95.7%)

We notice that the final preference model performs better than the initial preference model. We also observe an expected trend of the performance progressively improving from the initial preference model to the two intermediate preference models to the final preference model.

We observed similar trends in performance improvement for the preference models for lung and abdominal tumors. Thus the preference-assessment method leads to better preference models.

CONCLUSIONS

This paper describes a new preference-assessment method to overcome the difficulties in obtaining correct trade-off weights for attributes in multiattribute utility models. The preference-assessment method is based on the concept of knowledge maintenance for knowledge-based systems—the preference model is changed whenever the decision-theoretic system makes an incorrect recommendation. The interactive implementation ensures that the decision analyst is not needed after eliciting the initial preference model. The preference-assessment technique was implemented in the decision problem of evaluating competing three-dimensional radiation treatment plans.

An implicit assumption of the preference-assessment method is that the preference model will converge. We cannot predict the number of decision cases the decision maker will have to see before convergence. We were able to obtain convergence and demonstrate improvement in performance in three tumor sites. Lack of convergence can possibly mean that the multiattribute utility model being used is incorrect and needs to be changed.

The preference-assessment method is general and does not contain any domain information. Hence it can be used in other domains where the decision problem has properties such as multiple similar decision-making opportunities, necessity to make trade-offs, and no gold standard for decision making.

This research makes two key contributions. It describes a new preference-assessment method for assessing trade-offs for decision-theoretic systems which evaluate competing decision alternatives to fulfill multiple conflicting objectives using multiattribute utility theory. This includes the class of medical decision-making problems that involves the selection of the optimal therapy plan from competing therapy plans for a patient. It also describes a decision-theoretic plan-evaluation model and system for ranking competing 3D radiation treatment plans.

ACKNOWLEDGMENTS

Supported in part by National Cancer Institute Grant 5-U01-CA60245, and National Library of Medicine Grants 5-R29-LM05387 and U01-LM05845. We would like to thank Dr. Mary V. Graham, Dr. Henry K. Lee, and Dr. Jeff M. Michalski of the Division of Radiation Oncology for participating in the preference assessment experiments, and Dr. James A. Purdy of the Physics Section for his support and advice.

References

1. Miller PL, ed. Selected Topics in Medical Artificial Intelligence. New York: Springer-Verlag, 1988.
2. Miller RA. Medical diagnostic decision support systems—past, present, and future: A threaded bibliography and commentary. J Am Med Informatics Assoc 1994; 1: 8–27.
3. Heckerman DE, Horvitz EJ, Nathwani BN. Toward normative systems: Part I. The Pathfinder project. Meth Inform Med 1992; 31: 90–105.
4. Seiver A. A decision class analysis of critical care life-support decision-making. Int J Clin Monit Comput 1993; 10: 31–66.
5. Raiffa H. Decision Analysis: Introductory Lectures on Choice under Uncertainty. New York: Random House, 1968.
6. Keeney RL, Raiffa H. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. New York: John Wiley and Sons, 1976.
7. von Winterfeldt D, Edwards W. Decision Analysis and Behavioral Research. Cambridge, UK: Cambridge University Press, 1986.
8. Kassirer JP, Moskovitz AJ, Lau J, Pauker SG. Decision analysis: A progress report. Ann Intern Med 1987; 106: 275–91.
9. Hart A. Knowledge Acquisition for Expert Systems. New York: McGraw-Hill, 1992.
10. Musen MA. An overview of knowledge acquisition. In: David JM, Krivine JP, Simmons R, eds. Second Generation Expert Systems. Berlin: Springer-Verlag, 1993: 415–38.
11. Boose JH. A survey of knowledge acquisition techniques and tools. Knowledge Acquisition 1989; 1: 3–37.
12. Bachant J, McDermott J. R1 revisited: Four years in the trenches. AI Magazine 1984; 5(3): 21–32.
13. Farquhar PH. Utility assessment methods. Manage Sci 1984; 30: 1283–1300.
14. Dyer JS, Fishburn PC, Steuer RE, Wallenius J, Zionts S. Multiple criteria decision making, multiattribute utility theory: The next ten years. Manage Sci 1992; 38: 645–54.
15. Kahneman D, Slovic P, Tversky A, eds. Judgment under Uncertainty: Heuristics and Biases. Cambridge, UK: Cambridge University Press, 1982.
16. Stillwell WG, Seaver DA, Edwards E. A compar-

ison of weight approximation techniques in multiattribute utility decision making. Org Behav Human Decis Proc 1981; 28: 62–77.
17. Fischer GW. Utility models for multiple objective decisions: Do they accurately represent human preferences? Decis Sci 1979; 10: 451–79.
18. Borcherding K, Eppel T, von Winterfeldt D. Comparison of weighting judgments in multiattribute utility measurement. Manage Sci 1991; 37: 1603–19.
19. Barron FH. Selecting a best multiattribute alternative with partial information about attribute weights. Acta Psychologica 1992; 80: 91–103.
20. Delquié P. Inconsistent trade-offs between attributes: New evidence in preference assessment biases. Manage Sci 1993; 39: 1382–95.
21. Compton P, Edwards G, Kang B, et al. Ripple down rules: Turning knowledge acquisition into knowledge maintenance. Artif Intel Med 1992; 4: 463–75.
22. Edwards G, Compton P, Malor R, Srinivasan A, Lazarus L. PEIRS: A pathologist-maintained expert system for the interpretation of chemical pathology reports. Pathology 1993; 25: 27–34.
23. Farr BR, Shachter RD. Representation of preferences in decision-support systems. Comp Biomed Res 1992; 25: 324–35.
24. Rutledge GW, Thomsen GE, Farr BR, et al. The design and implementation of a ventilator-management advisor. Artif Intel Med 1993; 5: 67–82.
25. Perez CA, Brady LW, eds. Principles and Practice of Radiation Oncology. Second edition. Philadelphia: J. B. Lippincott, 1992.
26. Kutcher GJ, Leibel SA, Mohan R, et al. Advances in precision treatment: Some aspects of 3D conformal radiation therapy. Front Radiat Ther Oncol 1993; 27: 209–26.
27. Jain NL, Kahn MG. Objective evaluation of radiation treatment plans. In: Safran C, ed. Proceedings of the Seventeenth Annual Symposium on Computer Applications in Medical Care. New York: McGraw-Hill, 1993: 134–8.
28. Dyer JS, Sarin RK. Measurable multiattribute value functions. Oper Res 1979; 27: 810–22.
29. Purdy JA, Harms WB, Matthews JM, et al. Advances in 3-dimensional radiation treatment planning systems: Room-view display with real time interactivity. Int J Radiat Oncol Biol Phys 1993; 27: 933–44.